

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Chaoxin C. Qiu et al.

Filed: Concurrently Herewith

For: RESOURCE MANAGEMENT ARCHITECTURE FOR USE IN INFORMATION
STORAGE ENVIRONMENTS

Serial No.: Unknown

Prior Group Art Unit: Unknown

Examiner: Unknown

Atty. Dkt: SURG:160

NUMBER: EL917070800US

I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" service, postage prepaid, under 37 CFR 1.10 on the date indicated below and is address to: Assistant Commissioner of Patents, Washington, D.C. 20231.

Diane C. Potts
Signature

10-3-01
Date

Assistant Commissioner For Patents
Washington, D.C. 20231

Dear Sir:

PRELIMINARY AMENDMENT TO REDUCE FILING FEES

Please amend the application as follows.

In the specification:

The rewritten clean versions of all the specification changes are provided below.
Attached at the end of this paper is an Appendix providing an indication of the changes relative to the prior version of the specification, as now required by Rule 121.

Please replace the paragraph from page 1, lines 1-5 with the following:

097043 100304
103007 25402660

Patent Application for
**“RESOURCE MANAGEMENT ARCHITECTURE FOR USE IN INFORMATION
STORAGE ENVIRONMENTS”**

Inventors: Chaoxin C. Qiu, Umesh Gupta, Scott C. Johnson, Sarma Kolavasi, Theodore S.
Webb, Richard W. Yu, and Mark J. Conrad

Please replace the paragraph from page 1, line 7 to page 2, line 5 with the following:

This application is a continuation of Application Serial No. 09/947,869, which was filed September 6, 2001 and is entitled “SYSTEMS AND METHODS FOR RESOURCE MANAGEMENT IN INFORMATION STORAGE ENVIRONMENTS”, which in turn claims priority from co-pending United States Patent Application Serial Number 09/879,810 filed on June 12, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN INFORMATION MANAGEMENT ENVIRONMENTS,” and also claims priority from co-pending Provisional Application Serial No. 60/285,211 filed on April 20, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN A NETWORK ENVIRONMENT,” and also claims priority from co-pending Provisional Application Serial No. 60/291,073 filed on May 15, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN A NETWORK ENVIRONMENT,” the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending United States Patent Application Serial No. 09/797,198 filed on March 1, 2001 which is entitled “SYSTEMS AND METHODS FOR MANAGEMENT OF MEMORY,” and also claims priority from co-pending United States Patent Application Serial No. 09/797,201 filed on March 1, 2001 which is entitled “SYSTEMS AND METHODS FOR MANAGEMENT OF MEMORY IN INFORMATION DELIVERY ENVIRONMENTS,” and also claims priority from co-pending Provisional Application Serial No. 60/246,445 filed on November 7, 2000 which is entitled “SYSTEMS AND METHODS FOR PROVIDING EFFICIENT USE OF MEMORY FOR NETWORK SYSTEMS,” and also claims priority from co-pending Provisional Application Serial No. 60/246,359 filed on November 7, 2000 which is entitled “CACHING ALGORITHM

FOR MULTIMEDIA SERVERS,” the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending United States Patent Application Serial Number 09/797,200 filed on March 1, 2001 which is entitled “SYSTEMS AND METHODS FOR THE DETERMINISTIC MANAGEMENT OF INFORMATION” which itself claims priority from Provisional Application Serial No. 60/187,211 filed on March 3, 2000 which is entitled “SYSTEM AND APPARATUS FOR INCREASING FILE SERVER BANDWIDTH,” the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending Provisional Application Serial No. 60/246,401 filed on November 7, 2000 which is entitled “SYSTEM AND METHOD FOR THE DETERMINISTIC DELIVERY OF DATA AND SERVICES,” the disclosure of which is incorporated herein by reference.

Please replace the paragraph from page 49, lines 16-21 with the following:

In yet other embodiments, cycle time may be modified or limited based on a number of factors. For example cycle time may be limited or capped by limiting read-ahead buffer size, for example, using Resource Model Equations (17B), (18B) or (19B). Cycle time may also be limited or capped by placing a set limit on the maximal buffer size (*e.g.*, by placing a 2MB limit on the maximal buffer size in a case where system throughput does not increase, or does not increase significantly, with any increase in the buffer size beyond 2MB).

In the claims:

Please cancel claims 1-175 and add the following new claims:

176. An I/O resource management system capable of managing I/O resources in an information delivery environment, comprising:

an I/O resource model capable of modeling utilization of one or more of said I/O resources; and

an I/O resource manager in communication with said I/O resource model, said I/O resource manager being capable of managing one or more of said I/O resources based at least in part on said modeled utilization.

177. The system of claim 176, wherein said I/O resource model is capable of modeling utilization of one or more of said I/O resources based at least in part on one or more system I/O performance characteristics associated with said I/O resources.

178. The system of claim 177, wherein the value of one or more of said system I/O performance characteristics is estimated.

179. The system of claim 177, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with at least one of said I/O resource manager or said I/O resource model, said I/O resource monitor being capable of monitoring the value of one or more of said system I/O performance characteristics.

180. The system of claim 176, wherein said I/O resources comprise at least one of file system resources, storage system resources, or a combination thereof.

181. The system of claim 176, wherein said information delivery environment comprises delivery of continuous media data from an information management system in communication with said I/O resource management system; wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; and wherein said I/O resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data.

182. The system of claim 181, wherein said information management system comprises a storage system, said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

183. The system of claim 182, wherein said information delivery environment comprises delivery of continuous media data from said information management system to a network; and wherein said information management system comprises a content delivery system configured to be coupled to a network.

184. The system of claim 183, wherein said content delivery system is configured to be coupled to a network at an endpoint of said network.

185. The system of claim 182, wherein said I/O resource manager is capable of allocating one or more I/O resources between background system I/O activities and delivery of said continuous media data.

186. The system of claim 185, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with at least one of said I/O resource manager or said I/O resource model, said I/O resource monitor being capable of monitoring background system processing activity; and wherein said I/O resource manager is capable of allocating said one or more I/O resources between background system I/O activities and delivery of said continuous media data based at least in part on said monitored background system processing activity.

187. The system of claim 182, wherein said I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

188. The system of claim 181, wherein said I/O resource model comprises an analytical-based resource model.

189. The system of claim 181, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with said I/O resource model; and wherein said I/O resource model comprises a measurement-based resource model.

190. An I/O resource management system capable of managing I/O resources for delivery of continuous media data to a plurality of viewers from a storage system including at least one storage device or at least one partitioned group of storage devices, said system comprising:

an I/O resource monitor, said I/O resource monitor being capable of monitoring one or more system I/O performance characteristics associated with said I/O resources;

an I/O resource model in communication with said I/O resource monitor, said resource model being capable of modeling utilization of one or more of said I/O resources based at least in part on said one or more monitored system I/O performance characteristics; and

an I/O resource manager in communication with said I/O resource model, said I/O resource manager being capable of managing one or more of said I/O resources based at least in part on said modeled utilization.

191. The system of claim 190, wherein said storage system comprises part of a content delivery system configured to be coupled to a network.

192. The system of claim 191, wherein said content delivery system is configured to be coupled to said network at an endpoint of said network.

193. The system of claim 191, wherein said storage system includes at least two storage devices or at least two partitioned groups of storage devices for delivery of said continuous media data.

194. The system of claim 193, wherein said resource monitor is capable of monitoring a workload distribution across said at least two storage devices or at least two partitioned groups of storage devices; wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; and wherein said I/O resource model is capable of modeling said I/O capacity based at least in part on a workload distribution across said at least two storage devices or two or more partitioned groups of storage devices; and wherein said resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

195. The system of claim 194, wherein said buffer memory space comprises a part of an integrated cache/buffer memory of said storage system; wherein said I/O resource monitor is capable of monitoring a number of viewers that are reading data from said two or more storage devices or partitioned groups of storage devices out of the total number of viewers being served by said storage system; and wherein said I/O resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous

media data to said viewers reading data from said two or more storage devices or partitioned groups of storage devices.

196. The system of claim 194, wherein I/O resource manager is capable of allocating said I/O resources between background processing activities and delivery of said continuous media data.

197. The system of claim 194, wherein I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

198. The system of claim 194, wherein individual storage devices of said at least two storage devices or partitioned groups of storage devices comprise storage disk drives; and wherein said I/O resource model is capable of modeling utilization of one or more of said I/O resources based at least in part on one or more monitored system I/O performance characteristics associated with said I/O resources, said I/O system performance characteristics comprising at least one of seek and rotation latency, estimated transfer rate, or a combination thereof.

199. The system of claim 191, wherein said I/O resources comprise I/O capacity and buffer memory space of said storage system.

200. The system of claim 199, wherein said storage system comprises at least two storage devices or two partitioned groups of storage devices; and wherein said one or more monitored system I/O performance characteristics comprise one or more system I/O performance characteristics at least partially reflective of workload distribution across said at least two storage devices or said at least two partitioned groups of storage devices.

201. The system of claim 200, wherein said one or more monitored system I/O performance characteristics comprise at least one of maximal aggregate consumption rate for each of said at least two storage devices or partitioned groups of storage devices, maximal aggregate number of viewers for each of said at least two storage devices or partitioned groups of storage devices, or a combination thereof.

202. The system of claim 201, wherein said resource manager is capable of managing one or more of said I/O resources for delivery of said continuous media data to said plurality of viewers based at least in part on said modeled utilization.

203. The system of claim 202, wherein said resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

204. The system of claim 202, wherein said I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

205. The system of claim 191, wherein said I/O resource manager is capable of monitoring said one or more system I/O performance characteristics at the logical volume level.

206. The system of claim 205, wherein said I/O resource monitor is capable of monitoring said system I/O performance characteristics of said at least one storage device or at least one partitioned group of storage devices at the logical volume level.

207. The system of claim 205, wherein said I/O resource monitor is capable of monitoring at least one of maximal aggregate consumption rate for said at least one storage device or partitioned group of storage devices, maximal aggregate number of viewers for said at least one storage device or partitioned group of storage devices, or a combination thereof.

208. The system of claim 205, wherein said storage system includes at least two storage devices or at least two partitioned groups of storage devices.

209. The system of claim 208, wherein said I/O resource monitor is capable of monitoring a workload distribution across said at least two storage devices or at least two partitioned groups of storage devices.

210. An information delivery storage system, said storage system comprising:

a storage management processing engine that includes an I/O resource manager, a logical volume manager, and a monitoring agent; said I/O resource manager, said logical volume manager, and said monitoring agent being in communication; and

at least one storage device or group of storage devices coupled to said storage management processing engine;

wherein said information delivery storage system comprises part of an information management system configured to be coupled to a network.

211. The system of claim 210, wherein said storage management processing engine comprises one or more processing modules that are capable of performing at least one of I/O resource monitoring, I/O resource modeling, I/O resource management, or a combination thereof.

212. The system of claim 211, wherein said I/O resource manager comprises a storage system workload monitor.

213. The system of claim 212, wherein said monitoring agent is capable of monitoring number of outstanding I/O requests in at least one storage device or group of storage devices; and wherein said storage system workload monitor is capable of monitoring a number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices.

214. The system of claim 213, wherein said information management system comprises a content delivery system; wherein delivered information comprises continuous media data; and wherein said storage system includes two or more storage devices or two or more partitioned groups of storage devices for delivery of said continuous media data.

215. The system of claim 214, wherein said monitoring agent is capable of monitoring a number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices; and wherein said storage system workload monitor is capable of:

monitoring a number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by said at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices;

[illegible][illegible]

- [illegible]

[illegible][illegible]

wherein said estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices is determined based at least in part on said estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices, and estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices.

217. The system of claim 214, wherein said I/O resource manager has knowledge of the number of plex for each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices; wherein said monitoring agent is capable of monitoring a number of outstanding I/O requests for each said plex; and wherein said storage system workload monitor is capable of:

monitoring a number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices, monitoring the aggregated data consumption rates for said number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices;

determining an estimated total number of viewers for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored number of viewers for each logical volume;

determining an estimated aggregated data consumption rate for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored aggregated data consumption rates;

determining an estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part

on said estimated total number of viewers for each said plex and said monitored number of outstanding I/O requests for each said plex; and

determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for each said plex and said monitored number of outstanding I/O requests for each said plex.

218. The system of claim 217, wherein said monitoring agent is further capable of determining a maximal number of outstanding I/O requests for each said plex; and wherein said storage system workload monitor is capable of:

determining an estimated workload distribution across said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said monitored maximal number of outstanding I/O requests for each said plex;

determining an estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated total number of viewers for each said plex and said estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices; and

determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for each said plex and estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices.

219. The system of claim 217, wherein each of said storage devices comprise storage disk drives.

220. The system of claim 215, wherein said storage system workload monitor is capable of determining a maximal total number of viewers per storage device and a maximal aggregated data consumption rate storage device or per partitioned group of storage devices.

221. The system of claim 220, wherein said I/O resources comprise I/O capacity; and wherein said I/O resource manager is capable of modeling said I/O capacity based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device.

222. The system of claim 221, wherein said one or more I/O resources further comprise buffer memory space of said information management system; and wherein said I/O resource manager is capable of managing said I/O resources by balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices; wherein said balancing is based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

223. The system of claim 222, wherein said I/O resource manager is capable of performing I/O admission control, determining read-ahead size, or a combination thereof; wherein said performance of I/O admission control and determination of read-ahead size are based at least in part on said determined maximal total number of viewers per storage device or per partitioned

group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

224. The system of claim 223, wherein said I/O resource manager is capable of performing I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

225. The system of claim 223, wherein said resource manager is capable of determining read-ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

226. The system of claim 223, wherein said resource manager is capable of performing said I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and wherein said method further comprises determining read-ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

227. The system of claim 214, wherein said storage system workload monitor is capable of monitoring the following system I/O performance characteristics for each logical volume, for each plex within a logical volume, and for each storage device or partitioned group of storage devices within a plex: (1) total number of viewers, (2) aggregated data consumption rate, (3)

current weight of workload on a storage device in a plex, and (4) number of outstanding I/O requests for each storage device or partitioned group of storage devices.


228. The method of claim 214, wherein said storage system workload monitor is capable of determining a workload weight distribution for each of said storage devices or partitioned group of storage devices based at least in part on said monitored number of outstanding I/O requests for each storage device or partitioned group of storage devices.

REMARKS

Support for the new claims may be found in the claims as originally filed and throughout the Specification. No New Matter is added.

The examiner is invited to contact the undersigned at the phone number indicated below with any questions or comments, or to otherwise facilitate expeditious and compact prosecution of the application.

Respectfully submitted,



William W. Enders
Registration No. 41,735
Attorney for Applicant

O'KEEFE, EGAN & PETERMAN
1101 Capital of Texas Highway South
Building C, Suite 200
Austin, Texas 78746
(512) 347-1611
FAX: (512) 347-1615

APPENDIX
AMENDMENTS TO THE SPECIFICATION

Replacement paragraph for page 1, lines 1-5:

Patent Application for
“[SYSTEMS AND METHODS FOR RESOURCE MONITORING] RESOURCE
MANAGEMENT ARCHITECTURE FOR USE IN INFORMATION STORAGE
ENVIRONMENTS”

Inventors: Chaoxin C. Qiu, Umesh Gupta, Scott C. Johnson, Sarma Kolavasi, Theodore S.
Webb, Richard W. Yu, and Mark J. Conrad

Replacement paragraph for page 1, line 7 to page 2, line 5:

This application is a continuation of Application Serial No. 09/947,869, which was filed September 6, 2001 and is entitled “SYSTEMS AND METHODS FOR RESOURCE MANAGEMENT IN INFORMATION STORAGE ENVIRONMENTS”, which in turn [This application] claims priority from co-pending United States Patent Application Serial Number 09/879,810 filed on June 12, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN INFORMATION MANAGEMENT ENVIRONMENTS,” and also claims priority from co-pending Provisional Application Serial No. 60/285,211 filed on April 20, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN A NETWORK ENVIRONMENT,” and also claims priority from co-pending Provisional Application Serial No. 60/291,073 filed on May 15, 2001 which is entitled “SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN A NETWORK ENVIRONMENT,” the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending United States Patent Application Serial No. 09/797,198 filed on March 1, 2001 which is entitled “SYSTEMS AND METHODS FOR MANAGEMENT OF MEMORY,” and also claims priority from co-pending United States Patent Application Serial No. 09/797,201

filed on March 1, 2001 which is entitled "SYSTEMS AND METHODS FOR MANAGEMENT OF MEMORY IN INFORMATION DELIVERY ENVIRONMENTS," and also claims priority from co-pending Provisional Application Serial No. 60/246,445 filed on November 7, 2000 which is entitled "SYSTEMS AND METHODS FOR PROVIDING EFFICIENT USE OF MEMORY FOR NETWORK SYSTEMS," and also claims priority from co-pending Provisional Application Serial No. 60/246,359 filed on November 7, 2000 which is entitled "CACHING ALGORITHM FOR MULTIMEDIA SERVERS," the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending United States Patent Application Serial Number 09/797,200 filed on March 1, 2001 which is entitled "SYSTEMS AND METHODS FOR THE DETERMINISTIC MANAGEMENT OF INFORMATION" which itself claims priority from Provisional Application Serial No. 60/187,211 filed on March 3, 2000 which is entitled "SYSTEM AND APPARATUS FOR INCREASING FILE SERVER BANDWIDTH," the disclosures of each of the forgoing applications being incorporated herein by reference. This application also claims priority from co-pending Provisional Application Serial No. 60/246,401 filed on November 7, 2000 which is entitled "SYSTEM AND METHOD FOR THE DETERMINISTIC DELIVERY OF DATA AND SERVICES," the disclosure of which is incorporated herein by reference.

Replacement paragraph for page 49, lines 16-21:

In yet other embodiments, cycle time may be modified or limited based on a number of factors. For example cycle time may be limited or capped by limiting read-ahead buffer size, for example, using Resource Model Equations (17B), (18B) or (19B). Cycle time may also be limited or capped by placing a set limit on the maximal buffer size (*e.g.*, by placing a 2MB limit on the maximal buffer size in a case where system throughput does not increase, or does not increase significantly, with any increase in the buffer size beyond 2MB).